Microsoft®

Windows.net™

Server Family

Scale up
with
Windows® .Net Server

Jim Livingston
Lead Program Manager
Windows Datacenter Server
Microsoft Corporation
2/6/2002

Microsoft Windows.net Server Family

Defining Scalability

- **"Scalability": measure of increasing performance as computing resources are added to a system**
- **"Scaling up", or vertical scaling, refers to increasing server capacity by adding memory, processors, and I/O bandwidth to a single system**
- **"Scaling out", or horizontal scaling, refers to increasing server capacity by adding computer systems to grow overall server capacity**

Scaling up and scaling out are two deployment methods, and are not mutually exclusive

Scale-Up Factors

- When applications are architected for a single system
  - Most common examples are:
    - Databases
    - Data intensive applications
    - Applications with serial logic
    - Transaction processing applications
- Provides headroom to growing
- Operational savings
- Managing fewer physical systems is a priority for our customers

Windows Scale Architecture

Web/Presentation | **Business Logic** | **Data**

Microsoft Windows 2000 Server

Microsoft Application Center 2000

Microsoft Windows 2000 Advanced Server

Microsoft SQL Server 2000 Enterprise Edition

Microsoft Windows 2000 Datacenter Server

Scale up with Windows Server

- High-end Intel Servers
- Windows .Net Datacenter enhancements
- Stringent Hardware Test for reliability
  - Datacenter HCT – high end testing
- Windows IA64 support and status
- Multiple applications on single server
  - Server consolidation
- Resource Management

- Scale up server road map

High End Intel Servers

- OEMs delivering 32 processor systems
  - Windows support in Windows 2000 Datacenter
  - Unisys ES7000 (CMP) servers – 18 machines in Redmond
- NUMA based servers designs
  - More modular, flexible and easer to grow and maintain
  - Testing with first NUMA systems
- Partitioned systems
  - 8 static partitions supported in Windows 2000 server
  - Becoming common in server consolidation projects
  - Key feature for future servers
- IA32 machines (32 processors and 64 GB memory **maximum**)
  - 2 GB address space **maximum**
- IA64 Datacenter class machines in Redmond today
  - Testing with (16 processors \ 126 GB) and (32 processors \ 64 GB)
- Several major OEM’s will be shipping 256 processor servers by 2003

Inflection Point in Server Design

(Non Uniform Memory Access)

ccNUMA - Cache coherency means that the hardware is designed to allow each CPU to access all of the memory in the system as though it were directly attached to its own memory bus.

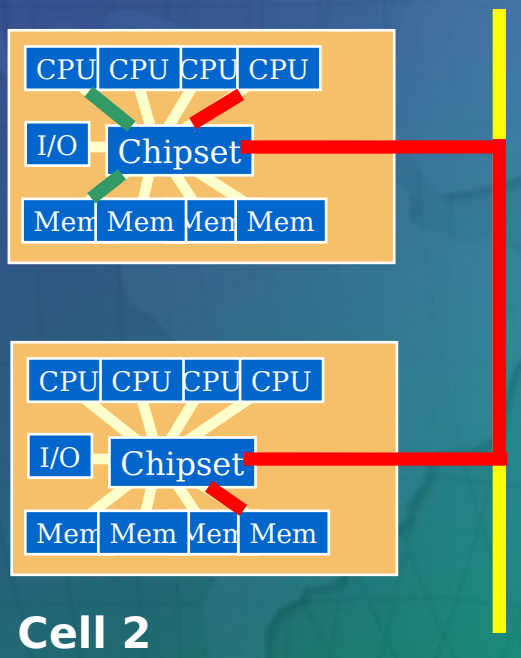

Local memory - access time to memory within the same cell is via the chipset in the cell and considered uniform

Remote memory - access time to memory in another cell, is via the local chipset, out through the system interconnect and through the chipset in the other cell. Typically longer.

Windows .NET Server recognizes a NUMA server and optimizes process scheduling and memory allocation to cells to deliver best performance.

Windows .Net Datacenter Server

- Large investment in Scale up performance
  - Improve SMP Scaling > 8 processors
  - SQL Server™, File, Print, IIS…
- Continued investment in Datacenter features
  - Provide IA64 and large memory support
    - OS Machine Check Architecture
  - NUMA enhancements
    - Applications \ processes kept on node
  - More Kernel performance improvements
  - RAS features
    - Memory mirroring
    - Support for hot add memory
  - Multi-Path IO
  - 8 Node Clustering
- Ensuring integration with largest systems
  - Lab dedicated to testing largest servers
  - Consolidation scenarios

Improvement in Scalability

- Kernel Improvements:
  - New and wider use of queued spinlocks
  - Process creation time reduced by 40%
  - Reduction in usage and hold time of key locks
    - Scheduling and Memory management
- File Server: Throughput 25% Higher on 4P
  - Reduction in context swaps and kernel improvements
  - ChkDsk: 60% to 6x faster than Windows 2000
  - IO Path and Scaling Improvements
- IIS: CGI currently 25% faster on 4P over Windows 2000
  - Process Creation and Heap Improvements
- Database: TPC-C 70% faster on 32P with kernel improvements
- Active Directory throughput improved from 2x to 3.9x on 8 way servers
- TPC-W 60% better through than Windows 2000

Application Benchmarks

Product	Windows/SQL Server Win Results	Rank
PeopleSoft CRM	**30,000 concurrent users**	**#1**
PeopleSoft eBill Payment	**191,694 payments per hour**	**#1**
Onyx	**32,000 concurrent users**	**#1**
Pivotal eRelationship	**20,000 concurrent users**	**#1**
SAP Retail	**3.165 million line items per hour**	**#1**
Great Plains	**2,400 concurrent users**	**#1**
SAP-SD Three Tier	**24,000 concurrent users**	**#2**

Results as of 1/7/02. Source: http://www.microsoft.com/sql/evaluation/compare/benchmar

Windows Scale-Up Progress

Windows and SQL Server have overtaken Sun in the SAP-SD benchmark using just half the number of processors

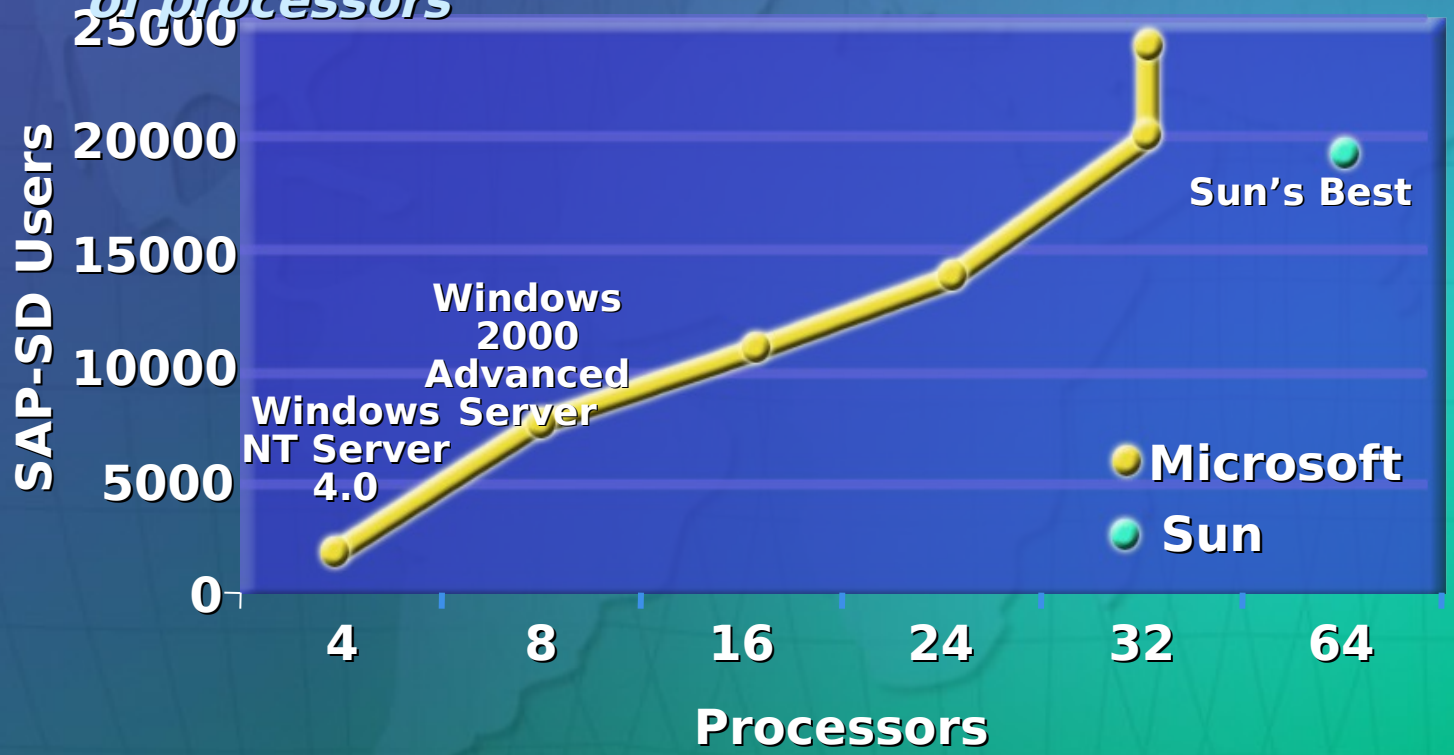

"SAP is convinced that SQL Server 2000 on Windows 2000 would now meet the needs of every R/3 customer around the world."

–Guenther Tolkmit, VP, SAP

Source: mySAP.com Standard Application Benchmarks, http://www.sap.com/solutions/technology/pdf/50020428.pdf

Microsoft Windows 2000 Datacenter Server

UNISYS We have a head for e-business.

Microsoft SQL Server 2000 Enterprise Edition

Industry Benchmark: TPC-C

Top TPC-C Non-Clustered Benchmark Performance

OS	Database	OEM	Available	$/perf	tpmC
Solaris 8	SymfoWARE	Fujitsu	2/28/2002	$28.58	455,818
HP UX 11i	Oracle 9i	HP	5/15/02	$21.24	389,434
Tru64 UNIX V5.1	Orcacle9i.9.0.1	Compaq	7/30/2001	$44.62	230,533
Solaris 8	SymfoWARE	Fujitsu/ICL	6/30/2001	$43.42	222,772
AIX 4.3.3	Oracle8 EE	IBM	4/13/2001	$34.18	220,807
AIX 4.3.3	Oracle 8i EE	Bull	5/28/2001	$34.67	220,807
HP UX 11.i 64-bit	Oracle8 EE	HP	5/1/2001	$43.25	197,024
Windows Datacenter Server Limited Edition	SQL Server 2000 Ent. Ed.	Unisys	3/10/2002	$21.33	165,219
IBM OS/400 V4 R5	IBM DB2	IBM	12/15/2000	$51.58	163,776
Tru64 UNIX V5.1	Oracle 8i EE	Compaq	2/2/2001	$52.88	155,179
IBM OS/400 V4 R5	IBM DB2	IBM	9/15/2000	$44.52	152,346
Windows Datacenter Server Limited Edition	SQL Server 2000 Ent. Ed.	Unisys	3/10/2002	$23.84	141,138

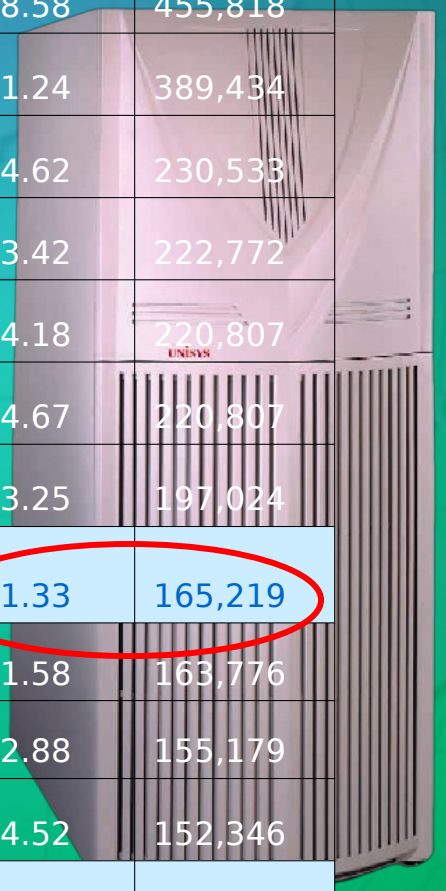

Results as of 1-8-02. Source: www.tpc.org

Datacenter Hardware Testing

- Stringent system validation required (HCT)
- 14-days Server stress test
- All key Server services stress to the maximum
  - Servers loaded past design points
  - Random workloads for (SQL, Exchange, File, Print, IIS...
  - Validation driven by 8 “load generators systems”
- 3-days re-test for “minor” configuration changes
  - Firmware
  - New adapters
  - Service Packs
- All “kernel” mode components must be loaded
- Configuration tool enables customers to check if their system is still validated
- Signed system configuration files downloadable from Datacenter HCL

ES7000s with Datacenter

78% had Zero Unplanned Downtime July – Dec 2001

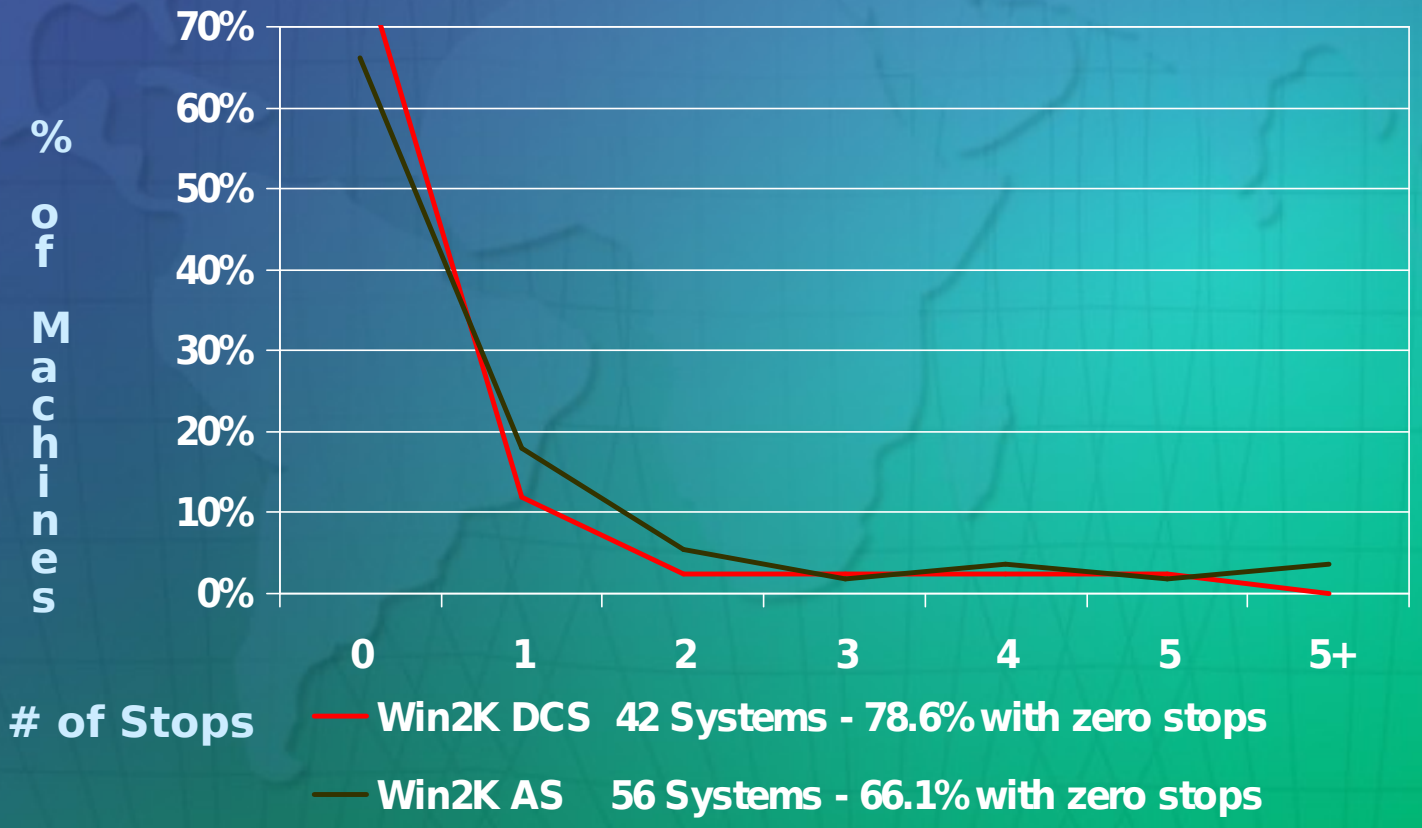

ISV Evangelism

- Additional resources and Datacenter class machines added to application compatibility lab
- Design reviews with ISV
  - Starting with the top 20 ...
- Scale up performance evaluations in Redmond
  - Windows Performance team
- Information on how to:
  - Scale their applications
  - Support consolidation
  - Provide 64 bit solutions
  - Clustered applications
  - Design for reliability
- Joint marketing activities to generate demand and promote products

Resource Management Strategy

- Partner with Aurema (ARMTech)
  - Provides solution to customers now
  - Testing completed in Redmond
  - ARMTech available to all Windows 2000 Datacenter customers
  - Product Support provided by Aurema
- Design and release our product with Windows .Net Server

Resource Management Scenarios

- Single or multiple LOB apps with other apps or services
  - **Example:** Fortis Health => Erisco + SQL Server 7 + SQL Server 2000 + Exchange 2000 + BizTalk™ Server + Commerce 2000 Server
- Multiple SQL Server instances
  - **Example:** MetLife => multiple departmental database servers consolidated as separate SQL Server instances running on single Datacenter machine
- Large Terminal Server system managing users
  - Managing resources used by users or groups of users
- Manage resource usage of individual IIS6 websites on a server
  - Single machine serving multiple websites
    - IIS6 allows separate processes can be assigned to each site
    - Command line matching feature of WRM matches policy to correct w3wp.exe process
- SQL Server, IIS and Exchange running on the same machine
  - Set policies for resource usage of each

Windows .Net Resource Management

- Customer Managed Resource allocation
  - Policy based allocation of system resources to applications on Windows Server
- Allows the customer to set allocation (CPU and memory) policies on applications
  - Select processes to be managed
  - Set resource consumption targets or limits
- Managed resources
  - CPU utilization (percent CPU)
  - Process working set size (physical resident pages)
  - Committed memory (pagefile usage)
- Apply policies on a date and time schedule
- Generate, store, view and export accounting records
- Automatically balance processes across nodes on NUMA machines

Windows Datacenter Roadmap

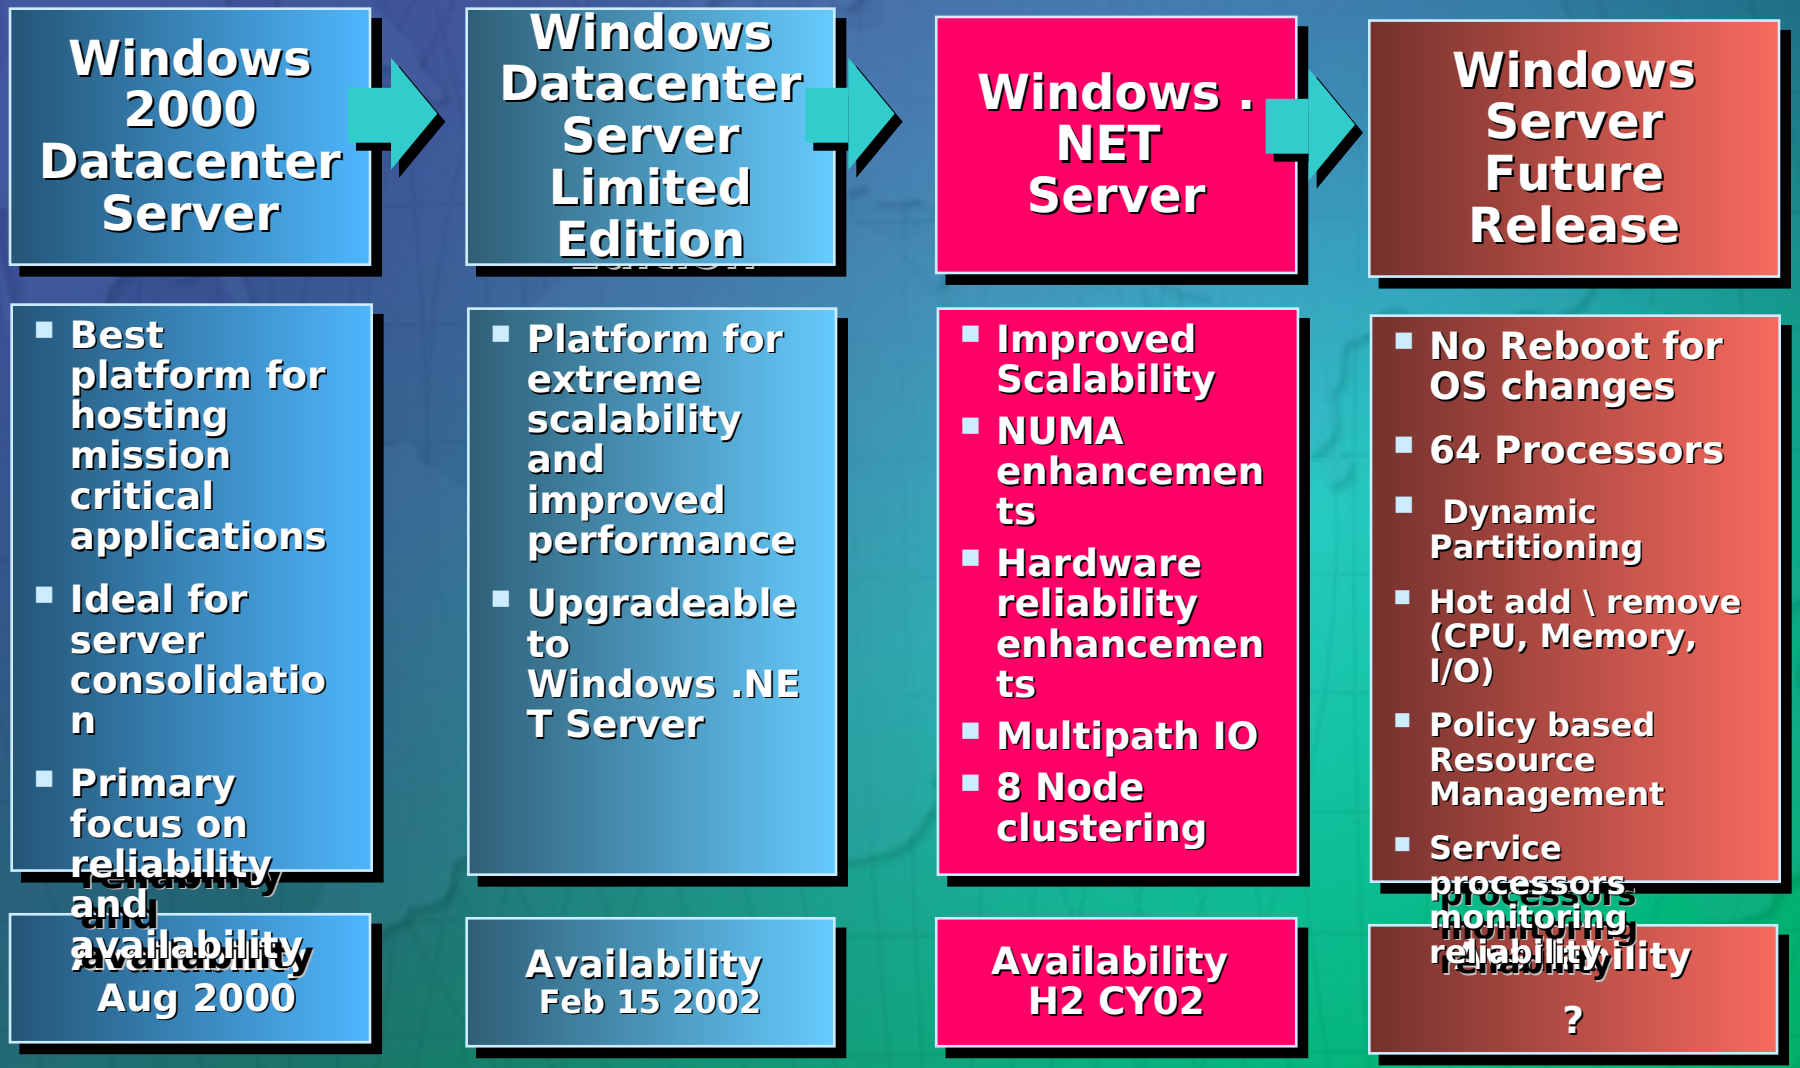

Datacenter Momentum

- Customers
  - Hundreds of Datacenter Servers sold and deployed
  - Licenses growing by 50% per quarter
- ISV
  - Hundreds of supported Applications
  - 23 Certified Applications and growing
- System Integrators
  - Major SIs are engaged
- Reliability
  - Achieving 99.999% availability
- Scalability
  - Evidence including SAP, TPC-C, TPC-W

Types of Server Consolidation

LOGICAL	GEOGRAPHIC
Centralizing Management	Centralizing Servers

HOMOGENEOUS	HETROGENEOUS
Single Application on a Single Box	Multiple Applications on a Single Box

Common Consolidation Scenarios

Scenario	Business Challenge	Windows Advantage
Database	**Too many duplicate servers running at low capacity.**	**Multiple databases, single box, 30K concurrent PeopleSoft users**
Messaging	**Messaging server farms growing to ensure scale**	**5000 concurrent users on a single Exchange server – Meta Group**
File & Print	**Sprawl of distributed servers without consistent management**	**W2K is 49% faster file server, 135% faster print server than Windows NT - Netbench.**

Server Consolidation

- Expand our engineering work on this scenario
  - Tested with Exchange, IIS, SQL …
  - Testing application maintenance solutions
  - WIST team simulating actual customer consolidation project
  - Providing guidelines, case studies …

Server Consolidation Customer

- Lay-Z-Boy
  - Largest manufacturer of reclining chairs
  - Consolidated onto 32 processors system
    - 12 SQL Servers plus Web servers and file and print servers
    - Centralize and consolidate to reduce management costs

“”To provide flexible and scalable services to our 10 different operating divisions, we want to deliver a platform for enterprise initiatives to drive down the total cost for all divisions, while still allowing each to maintain its individual uniqueness.”

Gary Clark, Director of Corporate IT Services, Lay-Z-Boy

LA-Z-BOY®

Microsoft® Windows® 2000 Datacenter Server

Server Consolidation Customer

- Abbey National
- Consolidate 25 Servers into one server
  - 16-processor Windows 2000 Datacenter Server partition:

2 different databases	SQL Server 2000 and Oracle 8.1.6
2 different business apps	Exchange 5.5, Home grown app
3 different environmental s/w	BMC Patrol, ESS, EMC Powerpath

  - 4-way UnixWare partition
  - 1.5TB database on EMC with 5m updates per day

"We've created an IT infrastructure with the reliability and scalability of a mainframe system and the flexibility of Windows technology."

Greg Walker, Manager Enterprise Group, Abbey National

ABBEY NATIONAL

Microsoft Windows 2000 Datacenter Server

Server Consolidation Customer

- Bank Labouchere
- E-trade leader in Europe, investment banking, corporate finance services
- Consolidated 250 Servers into 8 systems
  - IIS, Exchange, Oracle
  - Reduced the management burden
  - Reduce Total Cost of Ownership by 35%
  - Provide 5X headroom for changes in trading volumes
  - Provide 99.95% guaranteed business continuity, service availability to clients

64-bit Design Goals

- Deliver the best 64-bit operating system that:
  - Provides a solid foundation for the emergence of 64-bit Windows computing
  - Extends Windows to take full advantage of 64-bit computing
  - Enables Windows to support a new realm of applications and scenarios
  - Interoperates seamlessly with existing 32-bit deployments

Building 64-Bit Windows

- One Windows team - two hardware platforms
  - Dedication of entire Windows engineering force
  - Simultaneous release of both platforms a must
  - A showstopper on one blocks progress on both
- Driving Quality
  - Internal Test Focus:
    - Thousands of system-wide stress hours daily
    - Long haul server deployments, self-hosted desktops
  - Development Partners:
    - Partner evangelism, porting labs, SDK and DDK releases
  - Customers:
    - 3,000 64-bit beta site participants
    - Joint Development Program

64-Bit Windows Editions

- Windows XP
  - Released August 24th (with Windows XP)
- Windows Advanced Server Limited Edition
  - Refreshed in November
  - HP, Dell, Compaq, IBM ship servers
- Windows .Net Server release
  - Windows .NET Enterprise Server
  - Windows .NET Datacenter Server
- 32-Bit only:
  - Windows .NET Standard Server
  - Windows .NET Web Server
  - Windows Embedded

TCP\IP Connections

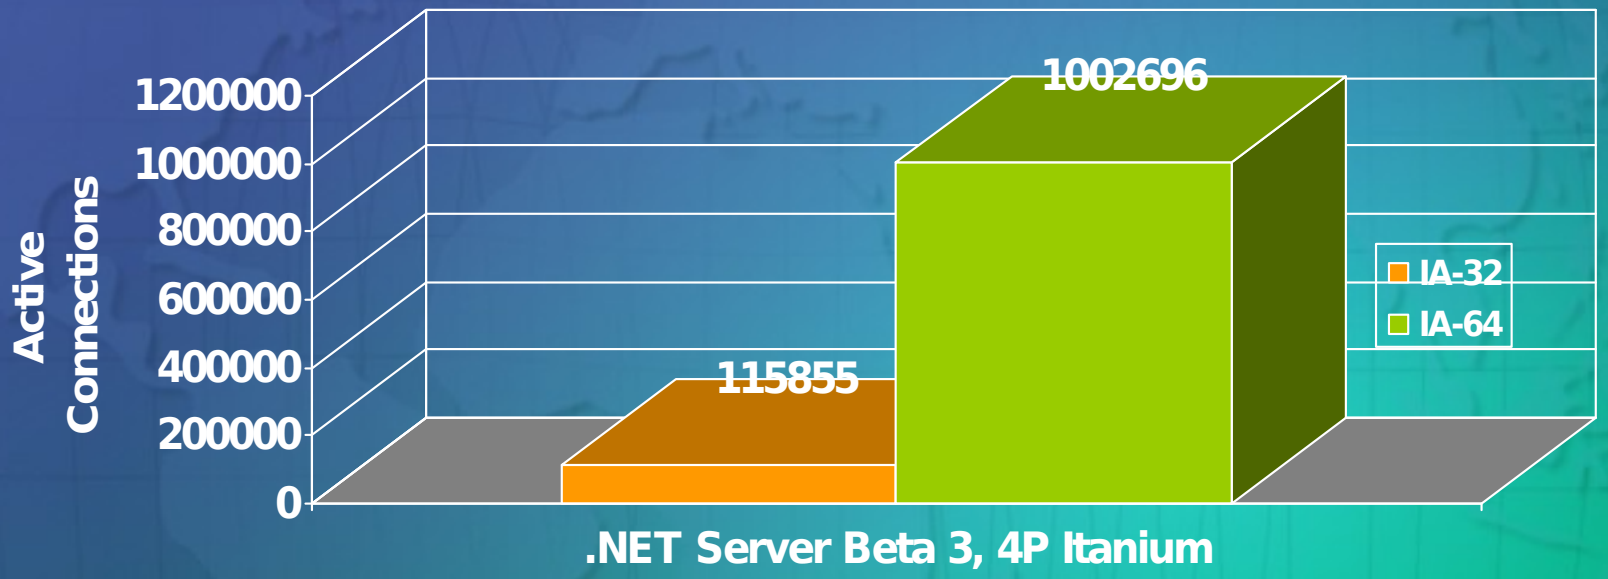

- Nonpaged Pool Usage:
  - 240 Mbyte on 32 bit (limited by VA)
  - 2.97 Gbyte on 64 bit (not even close to limit)

Process And Thread Creation

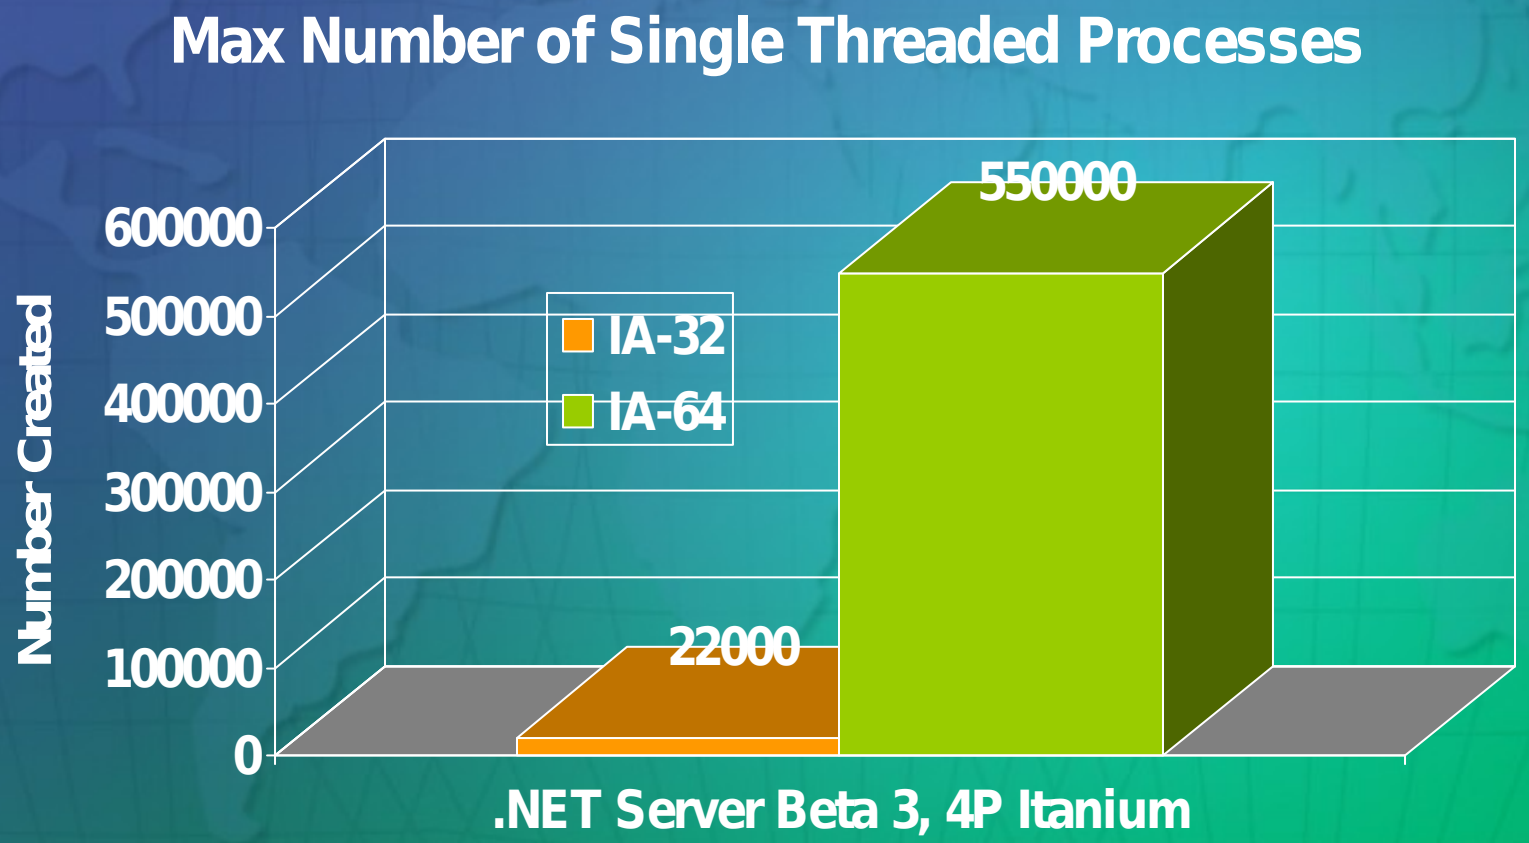

Machine Check Architecture

- Two components
  - Hardware: delivers certain hardware error events to the OS
  - Software: components that act on these event notices (e.g., logging, analysis, etc.)
  - Automatically corrects errors if possible
- Supported events
  - Hardware errors detected internally or on the Front Side Bus (FSB) of the CPU
  - Hardware errors reported to the CPU using the BERR or BINIT CPU pins

Not supported in initial releases

- Subsystems: DOS, 16-Bit, POSIX, OS/2
- .NET Framework, CLR, ASP.NET
- Jet RED
- Windows Media Services, WMP
- Fax Server/Client
- System Restore
- Remote Assistance
- Windows Product Activation
- Mobile features (sleep states, infrared)
- NetMeeting®

Memory Scorecard

Address Space	64-bit Windows	32-bit Windows
Virtual Memory	16 TB	4GB
Paging File	512 TB	16TB
Hyperspace	8 GB	4 MB
Paged Pool	128 GB	470 MB
Non-Paged Pool	128 GB	256 MB
System Cache	1 TB	1 GB

Scalable

- Platform for deploying large centralized database and data store back-ends
- Consolidate servers to reduce infrastructure size and complexity

Reliable

- Combines high-availability features with stringent platform validation and services
- Fastest path for problem resolution to minimize unplanned downtime

Better Investment

- Choice of high-end OEM suppliers while preserving application portfolio
- Powerful roadmap protects investments

Planning and Deployment Resources

- Information related to Microsoft® Scalability strategies and solutions:
  www.microsoft.com/servers/scalability

Microsoft®

Microsoft Windows.net Server Family